

Computing in the fractal cloud: modular generic solvers for SAT and Q-SAT variants

Denys Duchier, Jérôme Durand-Lose*, and Maxime Senot

LIFO, Université d'Orléans,
B.P. 6759, F-45067 ORLÉANS Cedex 2.

Abstract. Abstract geometrical computation can solve hard combinatorial problems efficiently: we showed previously how Q-SAT can be solved in bounded space and time using instance-specific signal machines and fractal parallelization. In this article, we propose an approach for constructing a particular *generic* machine for the same task. This machine deploies the Map/Reduce paradigm over a fractal structure. Moreover our approach is *modular*: the machine is constructed by combining modules. In this manner, we can easily create generic machines for solving satisfiability variants, such as SAT, #SAT, MAX-SAT.

Keywords. Abstract geometrical computation; Signal machine; Fractal; SAT; Massive parallelism; Model of computation.

1 Introduction

Since their first formulations in the seventies, problems of Boolean satisfiability have been studied extensively in the field of computational complexity. Indeed, the most important complexity classes can be characterized — in terms of reducibility and completeness — by such problems e.g. SAT for NP [Cook, 1971] and Q-SAT for PSPACE [Stockmeyer and Meyer, 1973]. As such, it is a natural challenge to consider how to solve these problems when investigating new computing machinery (quantum, NDA, membrane, hyperbolic spaces...) [Păun, 2001, Margenstern and Morita, 2001, Alhazov and Pérez-Jiménez, 2007].

This is the line of investigation that we have been following with *signal machines* [Durand-Lose, 2005], an abstract and geometrical model of computation. We showed previously how such machines were able to solve SAT [Duchier et al., 2010a] and Q-SAT [Duchier et al., 2010b] in bounded space and time. But in both cases, the machines were instance-specific *i.e.* depended on the formula whose satisfiability was to be determined. The primary contribution of the present paper is to exhibit a particular *generic* signal machine for the same task: it takes the instance formula as an input encoded (in polynomial-time by a Turing machine) in an initial configuration. We further improve our previous results by describing

* Corresponding author Jerome.Durand-Lose@univ-orleans.fr

a *modular* approach that allows us to easily construct generic machines for other variants of SAT, such as #SAT or MAX-SAT.

The model of signal machines, called *abstract geometrical computation*, involves two types of fundamental objects: dimensionless *particles* and *collision rules*. We use here one-dimensional machines: the space is the Euclidean real line, on which the particles move with a constant speed. Collision rules describe what happens when several particles collide. By representing the continuous time on a vertical axis, we obtain two-dimensional *space-time diagram*, in which the motion of the particles are materialized by segment lines called *signals*.

Signal machines can simulate Turing machines, and are thus Turing-universal [Durand-Lose, 2005]. They are also capable of analog computation by using the continuity of space and time to simulate analog models such as BSS's one [Durand-Lose, 2008, Blum et al., 1989] or computable analysis [Durand-Lose, 2009]. Other geometrical models of computation exist: colored universes [Jacopini and Sontacchi, 1990], geometric machines [Huckenbeck, 1989], piece-wise constant derivative systems [Bournez, 1997], optical machines [Naughton and Woods, 2001]... All these models, including signal machines, belong to a larger class of models of computation, called *unconventional*, which are more powerful than classical ones (Turing machines, RAM, λ -calculus ...). Among all these abstract models, the model of signal machines distinguishes itself by realistic assumptions respecting the major principles of physics — finite density of information, respect of causality and bounded speed of information — which are, in general, not respected all at the same time by other models. Nevertheless, signal machines remain an abstract model, with no a priori ambition to be physically realizable, and is studied for theoretical issues of computer sciences.

As signal machines take their origins in the world of cellular automata (as illustrated in Fig. 1), they can also be viewed as a massively parallel computational device. This is the approach proposed here: we put in place a fractal compute grid, then use the Map/Reduce paradigm to distribute the computations, then aggregate the results.

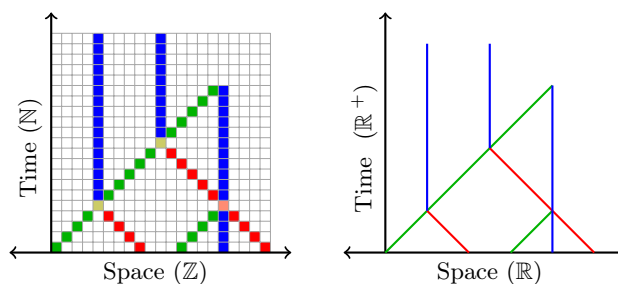


Fig. 1. From cellular automata to signal machines.

The Map/Reduce pattern, pioneered by Lisp, is now standard in functional programming: a function is applied to many inputs (map), then the results are aggregated (reduce). Google extended this pattern to allow its distributed com-

putation over a grid of possibly a thousand nodes [Dean et al., 2004]. The idea is to partition the input (petabytes of data) into chunks, and to process these chunks in parallel on the available nodes.

When solving combinatorial problems, we are also faced with massive inputs; namely, the exponential number of candidate solutions. Our approach is to distribute the candidates, and thus the computation, over an unbounded fractal grid. In this way, we adapt the map/reduce pattern for use over a grid with fractal geometry.

Our contribution in this paper is three fold: first, we show how Q-SAT can be solved in bounded space and time using a *generic machine*, where the input (the formula) is simply compiled into an initial configuration. This improves on our previous result where the machine itself depended on the formula. Second, we propose the first architecture for fractally distributed computing (the *fractal cloud*) and give a way to automatically shrink the data into this structure by means of a *lens device*. Third, we show how generic machines for many variants of SAT can be assembled by composing independent modules, which naturally emerged from the generalization of our previous family of machines into a single machine solving Q-SAT. Each module can be programmed and understood independently.

The paper is structured as follow. Signal machines are introduced in Section 2. Section 3 presents the fractal tree structure used to achieve massive parallelism and how general computations can be inserted in the tree. Section 4 details this implementation for a Q-SAT solver and Section 5 explains how some variants of satisfiability problems can be solved with the same approach. Complexities are discussed in Section 6 and conclusions and remarks are gathered in Section 7.

2 Definitions

Signal machines are an extension of cellular automata from discrete time and space to continuous time and space. Dimensionless signals/particles move along the real line and rules describe what happens when they collide.

Signals. Each *signal* is an instance of a *meta-signal*. The associated meta-signal defines its *velocity* and what happen when signals meet. Figure 2 presents a very simple space-time diagram. Time is increasing upwards and the meta-signals are indicated as labels on the signals.

Generally, we use over-line arrows to indicate the direction and speed of propagation of a meta-signal. For example, $\overrightarrow{\mathbf{a}}$ and $\overleftarrow{\mathbf{a}}$ denotes two different meta-signals; the first travels to the right at speed 1, while the other travels to the left at speed -3 . \mathbf{w} and \mathbf{a} are both stationary meta-signals.

Collision rules. When a set of signals collide, they are replaced by a new set of signals according to a matching collision rule. A rule has the form:

$$\sigma_1, \dots, \sigma_n \rightarrow \sigma'_1, \dots, \sigma'_p$$




Fig. 2. Geometrical algorithm for computing the middle

where all σ_i are meta-signals of distinct speeds as well as σ'_j (two signals cannot collide if they have the same speed and outgoing signals must have different speeds). A rule matches a set of colliding signals if its left-hand side is equal to the set of their meta-signals. By default, if there is no exactly matching rule for a collision, the behavior is defined to regenerate exactly the same meta-signals. In such a case, the collision is called *blank*. Collision rules can be deduced from space-time diagrams as on Fig. 2. They are also listed on the right of this figure.

Signal machine. A signal machine is defined by a set of meta-signals, a set of collision rules, and an initial configuration, i.e. a set of particles placed on the real line. The evolution of a signal machine can be represented geometrically as a *space-time diagram*: space is always represented horizontally, and time vertically, growing upwards. The geometrical algorithm displayed in Fig. 2 computes the middle: the new \mathbf{a} is located exactly halfway between the initial two \mathbf{w} .

3 Computing in the fractal cloud

Constructing the fractal: the fractal structure that interests us is based on the simple idea of computing the middle illustrated in Figure 2. We just indefinitely repeat this geometrical construction: once space has been halved, we recursively halve the two halves, and so on. This is illustrated in Figure 3(a), and can be generated by the following rules:¹

$$\mathbf{w}, \overleftarrow{\mathbf{a}} \rightarrow \mathbf{w}, \overrightarrow{\mathbf{a}} \qquad \overrightarrow{\mathbf{a}}, \overleftarrow{\mathbf{a}} \rightarrow \overleftarrow{\mathbf{a}}, \overleftarrow{\mathbf{a}}, \mathbf{a}, \overrightarrow{\mathbf{a}}, \overrightarrow{\mathbf{a}}$$

using $\{\mathbf{w}, \overrightarrow{\mathbf{a}}, \overleftarrow{\mathbf{a}}\} @ 0$ $\{\mathbf{w}\} @ 1$ as the initial configuration. This produces a stack of levels: each level is half the height of the previous one. As a consequence, the full fractal has width 1 and height 1.

¹ For brevity, we will always omit the rules which can be obtained from the others by symmetry. We refer to Appendix A for more details.

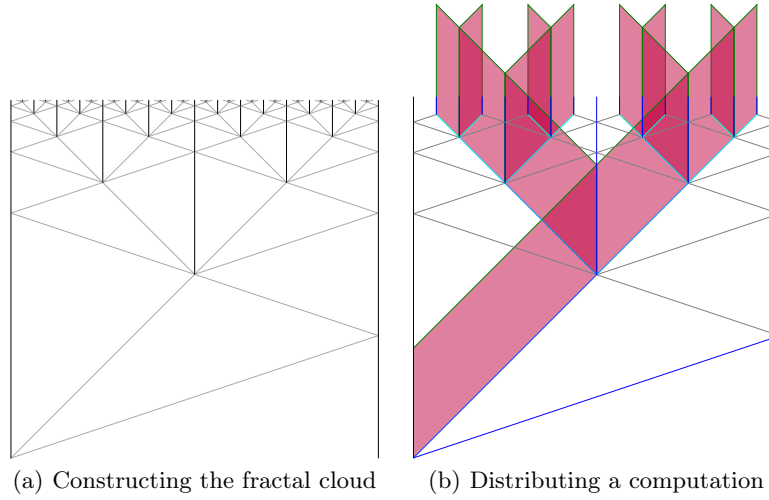


Fig. 3. Computing in the fractal cloud

Distributing a computation: the point of the fractal is to recursively halve space. At each point where space is halved, we position a stationary signal (a vertical line in the space-time diagram). We can use this structure, so that, at each halving point (stationary signal), we split the computation in two: send it to the left with half the data, and also to the right with the other half of the data.

The intuition is that the computation is represented by a *beam* of signals, and that stationary signals split this beam in two, resulting in one beam that goes through, and one beam that is reflected.

Unfortunately, a beam of constant width will not do: eventually it becomes too large for the height of the level. This can be clearly seen in Figure 3(b).

The lens device: the lens device narrows the beam by a factor of 2 at each level, thus automatically adjusting it to fit the fractal. It is implemented by the following meta-rule: *unless otherwise specified, any signal $\vec{\sigma}$ is accelerated by $\overleftarrow{\mathbf{a}}$ and decelerated and split by any stationary signal \mathbf{s} .*




Fig. 4. The lens device

Generic computing over the fractal cloud: with the lens device in effect, generic computations can take place over the fractal by propagating a beam from an initial configuration. We write $[(\vec{\sigma}_n \dots \vec{\sigma}_1) \text{spawn}]$ for an initial configuration with a sequence $\vec{\sigma}_n \dots \vec{\sigma}_1$ of signals. Geometrically, it can be seen easily that, in order for the beam to fit through the first level, the sequence $\vec{\sigma}_n \dots \vec{\sigma}_1$ must be placed in the interval $(-\frac{1}{4}, 0)$.

Stopping the fractal: For finite computations, we don't need the entire fractal. The $[\text{until}(n)]$ module can be inserted in the initial configuration to cut the fractal after n levels have been generated. $[\text{until}(n)] = \vec{z} \vec{\zeta}^{n-1}$

$$\begin{array}{lll} \vec{\zeta}, \overleftarrow{a} \rightarrow \overleftarrow{a}_o, \vec{\zeta} & \vec{\zeta}, a \rightarrow a_o & \vec{\zeta}, a_o \rightarrow \overleftarrow{\zeta}, a_o, \vec{\zeta} \\ \vec{z}, \overleftarrow{a} \rightarrow \overleftarrow{a}_o, \vec{z} & \vec{z}, \overleftarrow{a}_o \rightarrow \overleftarrow{a}, \vec{z} & \vec{z}, a_o \rightarrow \overleftarrow{z}, a, \vec{z} \\ \vec{z}, a \rightarrow a, \vec{z} & \vec{z}, \overrightarrow{a} \rightarrow \overrightarrow{a}_o & \overrightarrow{a}, \overleftarrow{a}_o \rightarrow \end{array}$$

The subbeam $\vec{\zeta}^{n-1}$ are the inhibitors for \vec{z} . One inhibitor is consumed at each level, after which \vec{z} takes effect and turns \overleftarrow{a} into \overleftarrow{a}_o , then crosses a and turns \overrightarrow{a} on the other side into \overrightarrow{a}_o . Finally the annihilation rule $\overrightarrow{a}, \overleftarrow{a}_o \rightarrow \emptyset$ brings the fractal to a stop. Thus, a computation $[(\vec{\sigma}_n \dots \vec{\sigma}_1 [\text{until}(n)]) \text{spawn}]$ uses only n levels. It can be seen geometrically that, for the collision of \vec{z} with \overrightarrow{a} to occur before the latter meets with \overleftarrow{a} , \vec{z} must initially be placed in $(-\frac{1}{6}, 0)$.

4 A modular Q-SAT solver

Q-SAT is the satisfiability problem for quantified Boolean formulae (QBF). A QBF is a closed formula of the form:

$$\phi = Q_1 x_1 Q_2 x_2 \dots Q_n x_n \quad \psi(x_1, x_2, \dots, x_n)$$

where $Q_i \in \{\exists, \forall\}$ and ψ is a quantifier-free formula of propositional logic. A recursive algorithm for solving Q-SAT is:

$$\begin{aligned} \text{qsat}(\exists x \phi) &= \text{qsat}(\phi[x \leftarrow \text{false}]) \vee \text{qsat}(\phi[x \leftarrow \text{true}]) \\ \text{qsat}(\forall x \phi) &= \text{qsat}(\phi[x \leftarrow \text{false}]) \wedge \text{qsat}(\phi[x \leftarrow \text{true}]) \\ \text{qsat}(\beta) &= \text{eval}(\beta) \end{aligned}$$

where β is a ground Boolean formula. This is exactly the structure of our construction: each quantified variable splits the computation in 2, $\text{qsat}(\phi[x \leftarrow \text{false}])$ is sent to the left and $\text{qsat}(\phi[x \leftarrow \text{true}])$ to the right, and subsequently the recursively computed results that come back are combined (with \vee for \exists and \wedge for \forall) to yield the result for the quantified formula. This process can be viewed as an instance of *Map/Reduce*, where the *Map* phase distributes the combinatorial exploration of all possible valuations across space using a binary decision tree,

and the *Reduce* phase collects the results and aggregates them using quantifier-appropriate Boolean operations. Our Q-SAT solver is modularly composed as follows:

`[(reduce:qsat($Q_1x_1 \dots Q_nx_n$)) [map:sat(ψ)] [decide(n)] [until($n+1$))] spawn]`

We describe the modules `decide`, `map:sat`, and `reduce:qsat` below.

4.1 Setting up the decision tree

For a QBF with n variables, we need 1 level per variable, and then at level $n+1$ we have a ground propositional formula that needs to be evaluated. Thus, the first module we insert is `[until($n+1$)]` to create $n+1$ levels. We then insert `[decide(n)]` because we want to use the first n levels as decision points for each variable.

$$[\text{decide}(n)] = \overrightarrow{\alpha}^n$$

$$\overrightarrow{\alpha}, \mathbf{a} \rightarrow \mathbf{x} \qquad \overrightarrow{\alpha}, \mathbf{x} \rightarrow \overleftarrow{\alpha}, \mathbf{x}, \overrightarrow{\alpha}$$

4.2 Compiling the formula

The intuition is that we want to compile the formula into a form of inverse polish notation to obtain executable code using postfix operators. At level $n+1$ all variables have been decided, and have become $\overrightarrow{\mathbf{t}}$ or $\overrightarrow{\mathbf{f}}$. The ground formula, regarded as an expression tree, can be executed bottom up to compute its truth value: the resulting signal for a subexpression is sent to interact with its parent operator.

The formula is represented by a beam of signals: each subformula is represented by a (contiguous) subbeam. A subformula that arrives at level $n+1$ starts evaluating when it hits the stationary \mathbf{a} . When its truth value has been computed, it is reflected so that it may eventually collide with the incoming signal of its parent connective.

Compilation. For binary connectives: one argument arrives first, it is evaluated, and its truth value is reflected toward the incoming connective; but, in order to reach it, it must cross the incoming beam for the other argument and not interact with the connectives contained therein. For this reason, with each subexpression, we associate a beam $\overrightarrow{\gamma}^k$ of inhibitors that prevents its resulting truth value to interact with the first k connectives that it crosses. We write $C[\psi]$ for the compilation of ψ into a contribution to the initial configuration, and $\|\psi\|$ for the number of occurrences of connectives in ψ .

$$\begin{aligned} C[\psi] &= C[\psi]^0 \\ C[\psi_1 \wedge \psi_2]^k &= \overrightarrow{\wedge} \overrightarrow{\gamma}^k C[\psi_1]^0 C[\psi_2]^{\|\psi_1\|} \\ C[\psi_1 \vee \psi_2]^k &= \overrightarrow{\vee} \overrightarrow{\gamma}^k C[\psi_1]^0 C[\psi_2]^{\|\psi_1\|} \\ C[\neg\psi]^k &= \overrightarrow{\neg} \overrightarrow{\gamma}^k C[\psi] \\ C[x_i]^k &= [\text{var}(x_i)] \overrightarrow{\gamma}^k \end{aligned}$$

Variables. We want variable x_i to be decided at level i . This can be achieved using $i - 1$ inhibitors. $[\text{var}(x_i)] = \overrightarrow{\mathbf{x}} \overrightarrow{\beta}^{i-1}$

$$\begin{array}{ll} \overrightarrow{\beta}, \mathbf{x} \rightarrow \mathbf{x}_o & \overrightarrow{\beta}, \mathbf{x}_o \rightarrow \overleftarrow{\beta}, \mathbf{x}_o, \overrightarrow{\beta} \\ \overrightarrow{\mathbf{x}}, \mathbf{x} \rightarrow \overleftarrow{\mathbf{f}}, \mathbf{x}, \overrightarrow{\mathbf{t}} & \overrightarrow{\mathbf{x}}, \mathbf{x}_o \rightarrow \overleftarrow{\mathbf{x}}, \mathbf{x}, \overrightarrow{\mathbf{x}} \end{array}$$

For variable x_i , the idea is to protect $\overrightarrow{\mathbf{x}}$ from being assigned into $\overleftarrow{\mathbf{f}}$ and $\overrightarrow{\mathbf{t}}$ until it reaches the i^{th} level. This is achieved with a stack of $i - 1$ signals $\overrightarrow{\beta}$: at each level, the first $\overrightarrow{\beta}$ turns the stationary signal \mathbf{x} into \mathbf{x}_o (the non-assigning version of \mathbf{x}), and dies. The following $\overrightarrow{\beta}$ are simply split, and so is $\overrightarrow{\mathbf{x}}$ but it additionally turns \mathbf{x}_o back into \mathbf{x} . After the first $i - 1$ levels, all the $\overrightarrow{\beta}$ have been consumed so that $\overrightarrow{\mathbf{x}}$ finally collides with \mathbf{x} and splits into $\overleftarrow{\mathbf{f}}$ going left and $\overrightarrow{\mathbf{t}}$ going right.

Evaluation. When hitting \mathbf{a} at level $n + 1$, $\overrightarrow{\mathbf{t}}$ is reflected as $\overleftarrow{\mathbf{T}}$, and $\overleftarrow{\mathbf{f}}$ as $\overleftarrow{\mathbf{F}}$: these are their *activated* versions which can interact with incoming connectives to compute the truth value of the formula according to the rules below (for \wedge ; other connectives are similar, cf Appendix A). See Fig. 5(a) for an example.

$$\begin{array}{lll} \overrightarrow{\mathbf{t}}, \mathbf{a} \rightarrow \overleftarrow{\mathbf{T}}, \mathbf{a} & \overrightarrow{\mathbf{f}}, \mathbf{a} \rightarrow \overleftarrow{\mathbf{F}}, \mathbf{a} & \overrightarrow{\gamma}, \mathbf{a} \rightarrow \overleftarrow{\gamma}_+, \mathbf{a} \\ \overrightarrow{\lambda}, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\lambda}_+ & \overrightarrow{\mathbf{f}}_+, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\mathbf{f}} & \overrightarrow{\lambda}_+, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\mathbf{t}} \\ \overrightarrow{\lambda}, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\mathbf{f}}_+ & \overrightarrow{\mathbf{f}}_+, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\mathbf{f}} & \overrightarrow{\lambda}_+, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\mathbf{f}} \\ \overrightarrow{\lambda}, \overleftarrow{\gamma}_+ \rightarrow \overrightarrow{\lambda}_o & \overrightarrow{\lambda}_o, \overleftarrow{\mathbf{T}} \rightarrow \overleftarrow{\mathbf{T}}, \overrightarrow{\lambda} & \overrightarrow{\lambda}_o, \overleftarrow{\mathbf{F}} \rightarrow \overleftarrow{\mathbf{F}}, \overrightarrow{\lambda} \end{array}$$

Storing the result. In order to make the result easily exploitable by the *Reduce* phase, we now store it as the stationary signal at level $n + 1$; it replaces \mathbf{a} . $[\text{store}] = \overrightarrow{\mathbf{s}}$

$$\overrightarrow{\mathbf{s}}, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\mathbf{T}} \quad \overrightarrow{\mathbf{s}}, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\mathbf{F}} \quad \overrightarrow{\mathbf{T}}, \mathbf{a} \rightarrow \mathbf{t} \quad \overrightarrow{\mathbf{F}}, \mathbf{a} \rightarrow \mathbf{f}$$

The complete *Map* phase is implemented by: $[\text{map: sat}(\psi)] = [\text{store}]C[\psi]$

4.3 Aggregating the results

As explained earlier, the results for an existentially (resp. universally) quantified variable must be combined using \vee (resp. \wedge).

Setting up the quantifiers. We turn the decision points of the first n levels into quantifier signals. Moreover, at each level, we must also take note of the direction in which the aggregated result must be sent. Thus \exists_L represents an existential quantifier that must send its result to the left.

$$[\text{reduce: qsat: init}(Q_1 x_1 \cdots Q_n x_n)] = \overrightarrow{Q_n} \cdots \overrightarrow{Q_1}$$

$$\mathbf{x}, \overleftarrow{\exists} \rightarrow \exists_R \quad \overrightarrow{\exists}, \mathbf{x} \rightarrow \exists_L \quad \mathbf{x}, \overleftarrow{\forall} \rightarrow \forall_R \quad \overrightarrow{\forall}, \mathbf{x} \rightarrow \forall_L$$




Fig. 5. Example $\exists x_1 \forall x_2 \forall x_3 (x_1 \wedge \neg x_2) \vee x_3$

Aggregating the results. Actual aggregation is initiated by \vec{c} and then executes according to the rules given in Fig 5(b). $[\text{reduce:qsat:exec}] = \vec{c}$

The complete Reduce phase is implemented by $[\text{reduce:qsat}(Q_1x_1 \cdots Q_nx_n)] = [\text{reduce:qsat:exec}][\text{reduce:qsat:init}(Q_1x_1 \cdots Q_nx_n)]$

5 Machines for SAT variants

Similar machines for variants of SAT can be obtained easily, typically by using different modules for the *Reduce* phase.

ENUM-SAT. returning all the satisfying assignments for a propositional formula ψ can be achieved easily by storing them as stationary beams.

$$[\text{reduce:allsat}(n)] = \vec{\mathbf{v}}[\text{var}(x_1)] \dots [\text{var}(x_n)] \vec{\mathbf{v}}$$

$$\begin{array}{llll} \vec{\mathbf{v}}, \mathbf{t} \rightarrow \overleftarrow{\mathbf{v}}_+, \mathbf{v} & \vec{\mathbf{v}}_+, \overleftarrow{\mathbf{t}} \rightarrow \mathbf{t}, \vec{\mathbf{v}}_+ & \vec{\mathbf{v}}, \mathbf{f} \rightarrow \overleftarrow{\mathbf{v}}_o & \vec{\mathbf{v}}_o, \overleftarrow{\mathbf{t}} \rightarrow \vec{\mathbf{v}}_o \\ \vec{\mathbf{v}}_+, \overleftarrow{\mathbf{v}} \rightarrow \mathbf{v} & \vec{\mathbf{v}}_+, \overleftarrow{\mathbf{f}} \rightarrow \mathbf{f}, \vec{\mathbf{v}}_+ & \vec{\mathbf{v}}_o, \overleftarrow{\mathbf{v}} \rightarrow & \vec{\mathbf{v}}_o, \overleftarrow{\mathbf{f}} \rightarrow \vec{\mathbf{v}}_o \end{array}$$

#SAT. counting the number of satisfying assignments for ψ can be achieved using signals for a binary adder. For lack of space, we cannot exhibit the rules here, but they can be found in Appendix B.

MAX-SAT. finding the maximum number of *clauses* that can be satisfied by an assignment. Here we must count the number of satisfied clauses rather than the number of satisfying assignments, and then stack a module for computing the max of two binary numbers.

6 Complexities

As mentioned in Sect. 1, we implement algorithms for satisfiability problems on signal machines in order to investigate the computational power of our abstract geometrical model of computation and to compare it to others. As we shall see, for such comparisons to be meaningful, the way complexity is measured must be adapted to the nature on the computing machine.




Fig. 6. The whole diagram.

Instead, we should regard our construction as a massively parallel computational device transforming inputs into outputs. The input is the initial configuration at the bottom of the diagram, and the output is the truth value signal coming out at the top, as seen in Fig. 6 for formula $\exists x_1 \forall x_2 \forall x_3 (x_1 \wedge \neg x_2) \vee x_3$. The transformation is performed in parallel by many threads: a thread here is an ascending path through the diagram from an input to the output, and the operations executed by the thread are the collisions occurring on this path.

Since signal machines can be regarded as the extension of cellular automata from discrete to continuous time and space, it might seem natural to measure the time (resp. space) complexity of a computation using the height (resp. width) of its space-time diagram. But, in our applications to SAT variants, these are bounded and independent of the formula: the *Map* phase is bounded by the fractal, and, by symmetry, so is the *Reduce* phase. Indeed, in general, by an appropriate scaling of the initial configuration, a finite computation could be made as small as desired. Thus, height and width are no longer pertinent measures of complexity.

Formally, we view a space-time diagram as a directed acyclic graph of collisions (vertices) and signals (arcs) oriented according to causality. Time complexity is then defined as the maximal size of a chain of collisions *i.e.* the length of the longest path, and space complexity as the maximal size of an anti-chain *i.e.* the size of the maximal set of signals pairwise un-related. This model-specific measure of time complexity is called *collisions depth*.

For the present construction, if s is the size of the formula and n the number of variables, space complexity is exponential: during evaluation, 2^n independent computations are executed in parallel, each one involving approximately s signals, so that the total space complexity is in $O(s \cdot 2^n)$.

Regarding the time complexity: for each subformula, the compilation process introduces a number of signals at most linear in s . Thus the number of signals in the initial configuration is $O(s^2)$. The primary contribution to the number of collisions along an ascending path comes, at each of the n levels, from the reflected beam crossing the incoming beam. Thus a thread involves $O(n \cdot s^2)$ collisions, making the collision depth cubic in the size of the formula instead of quadratic for our previous family of machines, giving us an idea of the price for genericity.

7 Conclusion

We showed in this paper that abstract geometrical computation can solve Q-SAT in bounded space and time by means of a single generic signal machine. This is achieved through massive parallelism enabled by a fractal construction that we call the *fractal cloud*. We adapted the Map/Reduce paradigm to this fractal cloud, and described a modular programming approach making it easy to assemble generic machines for SAT variants such as #SAT or MAX-SAT.

As we explained in Sect. 6, time and space are no longer appropriate measures of complexity for geometrical computations. This leads us to propose new definitions of complexity, specific to signal machines, and taking in account the parallelism of the model: time and space complexities are now defined respectively by the maximal sizes of a chain and an anti-chain, when the diagram is regarded as a directed acyclic graph. Time complexity thus defined is called *collision depth*. According to these new definitions, our construction has exponential space complexity and cubic time complexity.

Although the model is purely theoretical and has no ambition to be physically realizable, it is a significant and distinguishing aspect of signal machines that they solve satisfiability problems while adhering to major principles of modern physics — finite density and speed of information, causality — that are typically not considered by other unconventional models of computation. They do not, however, respect the quantization hypothesis, nor the uncertainty principle.

We are now furthering our research along two axes. First, the design and applications of other fractal structures for modular programming with fractal parallelism. Second, the investigation of computational complexity classes, both classical and model-specific for abstract geometrical computation.

Bibliography

- Artiom Alhazov and Mario J. Pérez-Jiménez. Uniform solution of Q-SAT using polarizationless active membranes. In J. Durand-Lose and M. Margenstern, editors, *5th International Conference on Machines, Computations, and Universality (MCU '07)*, number 4664 in LNCS, pages 122–133. Springer, 2007.
- L. Blum, M. Shub, and S. Smale. On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. *Bulletin of the American Mathematical Society*, 21(1):1–46, 1989.
- Olivier Bournez. Some bounds on the computational power of piecewise constant derivative systems. In *ICALP '97*, number 1256 in LNCS, pages 143–153, 1997.
- Stephen A. Cook. The complexity of theorem proving procedures. In *3rd Symposium on Theory of Computing (STOC '71)*, pages 151–158. ACM, 1971.
- Jeffrey Dean, Sanjay Ghemawat, and Google Inc. Mapreduce: simplified data processing on large clusters. In *Proceedings of the 6th Conference on Symposium on Operating Systems Design & Implementation (OSDI' 04)*. USENIX Association, 2004.
- Densy Duchier, Jérôme Durand-Lose, and Maxime Senot. Fractal parallelism: Solving SAT in bounded space and time. In O. Cheong, K-Y. Chwa, and K. Park, editors, *21st International Symposium on Algorithms and Computation (ISAAC '10)*, number 6506 in LNCS, pages 279–290. Springer, 2010a.
- Densy Duchier, Jérôme Durand-Lose, and Maxime Senot. Massively parallel automata in Euclidean space-time. In J. Beal, O. Michel, and A. Spicher, editors, *3rd International Spatial Computing Workshop (SCW '10)(extended abstracts of unpublished papers)*, 2010b. URL <http://www.spatial-computing.org>.
- Jérôme Durand-Lose. Abstract geometrical computation: Turing computing ability and undecidability. In B.S. Cooper, B. Löwe, and L. Torenvliet, editors, *New Computational Paradigms, 1st Conf. Computability in Europe (CiE '05)*, number 3526 in LNCS, pages 106–116. Springer, 2005.
- Jérôme Durand-Lose. Abstract geometrical computation with accumulations: Beyond the Blum, Shub and Smale model. In Arnold Beckmann, Costas Dimitracopoulos, and Benedikt Löwe, editors, *Logic and Theory of Algorithms, 4th Conf. Computability in Europe (CiE '08) (abstracts and extended abstracts of unpublished papers)*, pages 107–116. University of Athens, 2008.
- Jérôme Durand-Lose. Abstract geometrical computation and computable analysis. In J.F. Costa and N. Dershowitz, editors, *International Conference on Unconventional Computation 2009 (UC '09)*, number 5715 in LNCS, pages 158–167. Springer, 2009.
- Ulrich Haken. Euclidean geometry in terms of automata theory. *Theoret. Comp. Sci.*, 68(1):71–87, 1989.
- Giuseppe Jacopini and Giovanna Sontacchi. Reversible parallel computation: an evolving space-model. *Theoret. Comp. Sci.*, 73(1):1–46, 1990.
- Maurice Margenstern and Kenichi Morita. NP problems are tractable in the space of cellular automata in the hyperbolic plane. *Theor. Comp. Sci.*, 259(1-2):99–128, 2001.
- Thomas J. Naughton and Damien Woods. On the computational power of a continuous-space optical model of computation. In M. Margenstern, editor, *Machines, Computations, and Universality (MCU '01)*, number 2055 in LNCS, pages 288–299, 2001.
- Gheorghe Păun. P-systems with active membranes: Attacking NP-Complete problems. *Journal of Automata, Languages and Combinatorics*, 6(1):75–90, 2001.
- Larry J. Stockmeyer and Albert R. Meyer. Word problems requiring exponential time. In *5th ACM Symposium on Theory of Computing (STOC '73)*, volume 16, pages 1–9, 1973.

A Details of the modules

In the following, we generally define the collision rules only for one side, as the corresponding rules for the other side can be deduced by symetry. We give here the rules for all modules: first the modules for building and stopping the fractal tree, then the modules for setting and using the tree for satisfiability problems and finally the modules specific to each variants of SAT.

All the diagrams of the paper were generated by Durand-Lose's software, implemented in Java, and corresponds to a run of our Q-SAT solver for the formula:

$$\phi = \exists x_1 \forall x_2 \forall x_3 (x_1 \wedge \neg x_2) \vee x_3 \ .$$

A.1 The fractal cloud

Constructing the fractal: $[\text{start}] = \vec{a}, \overleftarrow{a}$

The following rules on the left correspond to the bounce of \vec{a} and \overleftarrow{a} on the wall w . Rules on the right are the step of induction for starting the next level: the initial signals \vec{a} and \overleftarrow{a} are duplicated on the right and the left, and the stationary signal a is created exactly at the middle of the previous stage. The result is given by Fig. 7.

$$\begin{array}{ll} w, \overleftarrow{a} \rightarrow w, \vec{a} & \vec{a}, \overleftarrow{a} \rightarrow \overleftarrow{a}, \overleftarrow{a}, a, \vec{a}, \vec{a} \\ \vec{a}, w \rightarrow \overleftarrow{a}, w & \vec{a}, \overleftarrow{a} \rightarrow \overleftarrow{a}, \overleftarrow{a}, a, \vec{a}, \vec{a} \end{array}$$




Fig. 7. The fractal tree.

Stopping the fractal: $[\text{until}(n+1)] = \vec{z} \vec{\zeta}^n$

The role of this module is to stop the fractal after $(n+1)$ levels — n levels for assigning the n variables and 1 level for the evaluation of the ground formula. For this, we use a stack of n signals $\vec{\zeta}$ and one signal \vec{z} . The $\vec{\zeta}$ signals are used both as a counter, one signal being killed at each level, and as inhibitors to the effect of \vec{z} . After

Case of $\vec{\alpha}$: the first $\vec{\alpha}$ is stopped by the stationary signal \mathbf{a} , which becomes \mathbf{x} and decelerates and splits the next coming $\vec{\alpha}$.

$$\vec{\alpha}, \mathbf{a} \rightarrow \mathbf{x} \qquad \vec{\alpha}, \mathbf{x} \rightarrow \overleftarrow{\alpha}, \mathbf{x}, \vec{\alpha}$$

Case of quantifiers signals: the first quantifier signal, $\vec{\forall}$ or $\vec{\exists}$, colliding with a stationary \mathbf{x} is stopped and turns \mathbf{x} into \forall_D or \exists_D ($D \in \{R, L\}$). This is achieved by the rules:

$$\mathbf{x}, \vec{\exists} \rightarrow \exists_R \qquad \vec{\exists}, \mathbf{x} \rightarrow \exists_L \qquad \mathbf{x}, \vec{\forall} \rightarrow \forall_R \qquad \vec{\forall}, \mathbf{x} \rightarrow \forall_L$$

The next quantifier signals are decelerated and split by the new stationary signal \forall_D or \exists_D . In the following rules, we have $D \in \{R, L\}$:

$$\begin{array}{ll} \vec{\exists}, \exists_D \rightarrow \overleftarrow{\exists}, \exists_D, \vec{\exists} & \vec{\exists}, \forall_D \rightarrow \overleftarrow{\exists}, \forall_D, \vec{\exists} \\ \vec{\forall}, \exists_D \rightarrow \overleftarrow{\forall}, \exists_D, \vec{\forall} & \vec{\forall}, \forall_D \rightarrow \overleftarrow{\forall}, \forall_D, \vec{\forall} \end{array}$$

A.2 The tree for satisfiability problems

Activation of the decision tree: $[\text{init}(n)] = \vec{\alpha}^n$

$$\vec{\alpha}, \mathbf{a} \rightarrow \mathbf{x} \qquad \vec{\alpha}, \mathbf{x} \rightarrow \overleftarrow{\alpha}, \mathbf{x}, \vec{\alpha}$$

Representation of a variable: $[\text{var}(x_i)] = \vec{\mathbf{x}} \vec{\beta}^{i-1}$

As we explained in the paper, the variable x_i is represented by a stack of i signals: one signal $\vec{\mathbf{x}}$ and $i-1$ signals $\vec{\beta}$. The role of the signals $\vec{\beta}$ is to protect $\vec{\mathbf{x}}$ from being assigned before the i^{th} level. The first signal $\vec{\beta}$ of the stack is stopped at the next split, so that $i-1$ signals have disappeared just before the i^{th} stage. This is illustrated in Fig. 8 and modelised by the following rules:

$$\begin{array}{ll} \vec{\beta}, \mathbf{x} \rightarrow \mathbf{x}_o & \vec{\beta}, \mathbf{x}_o \rightarrow \overleftarrow{\beta}, \mathbf{x}_o, \vec{\beta} \\ \vec{\mathbf{x}}, \mathbf{x} \rightarrow \overleftarrow{\mathbf{f}}, \mathbf{x}, \vec{\mathbf{t}} & \vec{\mathbf{x}}, \mathbf{x}_o \rightarrow \overleftarrow{\mathbf{x}}, \mathbf{x}, \vec{\mathbf{x}} \end{array}$$

Compilation of the formula: We propose here a recursive algorithm which takes as input an unquantified formula — a SAT-formula — and outputs the part of the initial configuration corresponding to the formula. In the following schemes of compilation, $\|\phi\|$ designates the number of occurrences of Boolean connectives in formula ϕ .

$$\begin{aligned} C[\phi] &= C[\phi]^0 \\ C[\phi_1 \wedge \phi_2]^k &= \vec{\wedge} \vec{\gamma}^k C[\phi_1]^0 C[\phi_2]^{\|\phi_1\|} \\ C[\phi_1 \vee \phi_2]^k &= \vec{\vee} \vec{\gamma}^k C[\phi_1]^0 C[\phi_2]^{\|\phi_1\|} \\ C[\neg\phi]^k &= \vec{\neg} \vec{\gamma}^k C[\phi] \\ C[x_i]^k &= \vec{\mathbf{x}} \vec{\beta}^{i-1} \vec{\gamma}^k \end{aligned}$$

Evaluation: The rules for the evaluation follow the classical Boolean operations. We explained earlier that some inhibiting signals — the $\overrightarrow{\gamma}$ — are needed to allow the result of the first evaluated argument of a binary connective to traverse the beam of the other, as yet unevaluated, argument without reacting with the connectives contained therein, and only interact with its actual syntactical parent connective. Figure 9 displays the evaluation of the formula $\phi = \exists x_1 \forall x_2 \forall x_3 (x_1 \wedge \neg x_2) \vee x_3$ for the case $x_1 = x_2 = x_3 = \mathbf{t}$.

$$\begin{array}{ccc}
\overrightarrow{\mathbf{t}}, \mathbf{a} \rightarrow \overleftarrow{\mathbf{t}}, \mathbf{a} & \overrightarrow{\mathbf{f}}, \mathbf{a} \rightarrow \overleftarrow{\mathbf{f}}, \mathbf{a} & \overrightarrow{\gamma}, \mathbf{a} \rightarrow \overleftarrow{\gamma}_+, \mathbf{a} \\
\\
\overrightarrow{\lambda}, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\lambda}_+ & \overrightarrow{\mathbf{f}}_+, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\mathbf{f}} & \overrightarrow{\lambda}_+, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\mathbf{t}} \\
\overrightarrow{\lambda}, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\mathbf{f}}_+ & \overrightarrow{\mathbf{f}}_+, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\mathbf{f}} & \overrightarrow{\lambda}_+, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\mathbf{f}} \\
\\
\overrightarrow{\nabla}, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\mathbf{t}}_+ & \overrightarrow{\mathbf{t}}_+, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\mathbf{t}} & \overrightarrow{\nabla}_+, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\mathbf{t}} \\
\overrightarrow{\nabla}, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\nabla}_+ & \overrightarrow{\mathbf{t}}_+, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\mathbf{t}} & \overrightarrow{\nabla}_+, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\mathbf{f}} \\
\\
\overrightarrow{\Rightarrow}, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\mathbf{f}} & & \\
\overrightarrow{\Rightarrow}, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\mathbf{t}} & & \\
\\
\overrightarrow{\lambda}, \overleftarrow{\gamma}_+ \rightarrow \overrightarrow{\lambda}_\circ & \overrightarrow{\nabla}, \overleftarrow{\gamma}_+ \rightarrow \overrightarrow{\nabla}_\circ & \overrightarrow{\Rightarrow}, \overleftarrow{\gamma}_+ \rightarrow \overrightarrow{\Rightarrow}_\circ \\
\\
\overrightarrow{\lambda}_\circ, \overleftarrow{\mathbf{T}} \rightarrow \overleftarrow{\mathbf{T}}, \overrightarrow{\lambda} & \overrightarrow{\nabla}_\circ, \overleftarrow{\mathbf{T}} \rightarrow \overleftarrow{\mathbf{T}}, \overrightarrow{\nabla} & \overrightarrow{\Rightarrow}_\circ, \overleftarrow{\mathbf{T}} \rightarrow \overleftarrow{\mathbf{T}}, \overrightarrow{\Rightarrow} \\
\overrightarrow{\lambda}_\circ, \overleftarrow{\mathbf{F}} \rightarrow \overleftarrow{\mathbf{F}}, \overrightarrow{\lambda} & \overrightarrow{\nabla}_\circ, \overleftarrow{\mathbf{F}} \rightarrow \overleftarrow{\mathbf{F}}, \overrightarrow{\nabla} & \overrightarrow{\Rightarrow}_\circ, \overleftarrow{\mathbf{F}} \rightarrow \overleftarrow{\mathbf{F}}, \overrightarrow{\Rightarrow}
\end{array}$$

Storing the results: $[\text{store}] = \overrightarrow{\mathbf{s}}$

$$\begin{array}{ccc}
\overrightarrow{\mathbf{s}}, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\mathbf{T}} & \overrightarrow{\mathbf{T}}, \mathbf{a} \rightarrow \mathbf{t} \\
\overrightarrow{\mathbf{s}}, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\mathbf{F}} & \overrightarrow{\mathbf{F}}, \mathbf{a} \rightarrow \mathbf{f}
\end{array}$$

B Modularity and satisfiability variants

B.1 Q-SAT

To proceed to the aggregation process, we join the results coming from right and left two-by-two. This is done with a stationary signal indicating the type of operation to perform — a conjunction for \forall and a disjunction for \exists — and the direction of the resulting signal — left or right. The whole process is displayed in Fig. 10.

Setting up the reduce stage:

$$\begin{array}{l}
[\text{reduce:qsat:init}(Q_1 x_1 \cdots Q_n x_n)] = \overrightarrow{Q_n} \dots \overrightarrow{Q_1} \\
\\
\mathbf{x}, \overleftarrow{\exists} \rightarrow \exists_R \qquad \qquad \qquad \overrightarrow{\exists}, \mathbf{x} \rightarrow \exists_L \\
\mathbf{x}, \overleftarrow{\forall} \rightarrow \forall_R \qquad \qquad \qquad \overrightarrow{\forall}, \mathbf{x} \rightarrow \forall_L
\end{array}$$




Fig. 9. Evaluation for $x_1 = x_2 = x_3 = \mathbf{t}$ in $\exists x_1 \forall x_2 \forall x_3 (x_1 \wedge \neg x_2) \vee x_3$.




Fig. 10. Aggregation process.

Executing the reduce stage: $[\text{reduce:qsat:exec}] = \vec{c}$

Initiation:

$$\vec{c}, t \rightarrow \overleftarrow{t}$$

$$\vec{c}, f \rightarrow \overleftarrow{f}$$

Performing the disjunction:

$$\begin{array}{cccc} \vec{t}, \exists_L, \overleftarrow{t} \rightarrow \overleftarrow{t} & \vec{t}, \exists_L, \overleftarrow{f} \rightarrow \overleftarrow{t} & \vec{f}, \exists_L, \overleftarrow{t} \rightarrow \overleftarrow{t} & \vec{f}, \exists_L, \overleftarrow{f} \rightarrow \overleftarrow{f} \\ \vec{t}, \exists_R, \overleftarrow{t} \rightarrow \vec{t} & \vec{t}, \exists_R, \overleftarrow{f} \rightarrow \vec{t} & \vec{f}, \exists_R, \overleftarrow{t} \rightarrow \vec{t} & \vec{f}, \exists_R, \overleftarrow{f} \rightarrow \vec{f} \end{array}$$

Performing the conjunction:

$$\begin{array}{cccc} \vec{t}, \forall_L, \overleftarrow{t} \rightarrow \overleftarrow{t} & \vec{t}, \forall_L, \overleftarrow{f} \rightarrow \overleftarrow{f} & \vec{f}, \forall_L, \overleftarrow{t} \rightarrow \overleftarrow{f} & \vec{f}, \forall_L, \overleftarrow{f} \rightarrow \overleftarrow{f} \\ \vec{t}, \forall_R, \overleftarrow{t} \rightarrow \vec{t} & \vec{t}, \forall_R, \overleftarrow{f} \rightarrow \vec{f} & \vec{f}, \forall_R, \overleftarrow{t} \rightarrow \vec{f} & \vec{f}, \forall_R, \overleftarrow{f} \rightarrow \vec{f} \end{array}$$

Putting all the modules together, we obtain for the running example the initial configuration shown by Fig. 11.

The global construction is displayed by Fig. 12



Fig. 11. Initial configuration.




Fig. 12. The whole diagram.

B.2 #SAT

#SAT is the problem of counting the number of solutions of SAT. We recall that this problem is complete for the class #P *i.e.* the class of NP-problems for which their solutions can be counted in polynomial time.

To solve #SAT, as a SAT-formula is a special instance of a Q-SAT formula in which all quantifiers are existential, we can use our Q-SAT solver and we add a special module for counting the truth evaluation of the formula during the aggregation process. The counting is performed by a binary adder.

Setting up the reduce stage: $[\text{reduce}:\#\text{sat}:\text{init}] = \overrightarrow{+}^n$

$$\overrightarrow{+}, \mathbf{x} \rightarrow +_L^0 \qquad \mathbf{x}, \overleftarrow{+} \rightarrow +_R^0$$

Executing the reduce stage: $[\text{reduce}:\#\text{sat}:\text{exec}] = \overrightarrow{\epsilon}_o \overrightarrow{\delta}_o \overrightarrow{0}_o$

$$\begin{array}{lll} \overrightarrow{\delta}_o, \mathbf{t} \rightarrow \overleftarrow{\delta} & \overrightarrow{\epsilon}_o, \mathbf{t} \rightarrow \overleftarrow{\epsilon} & \overrightarrow{0}_o, \mathbf{t} \rightarrow \overleftarrow{1} \\ \overrightarrow{\delta}_o, \mathbf{f} \rightarrow \overleftarrow{\delta} & \overrightarrow{\epsilon}_o, \mathbf{f} \rightarrow \overleftarrow{\epsilon} & \overrightarrow{0}_o, \mathbf{f} \rightarrow \overleftarrow{0} \end{array}$$

Rules for the binary adder (for the R subscript, the rules for the L are similar, but with the output going to the left):

$$\begin{array}{ll}
\overrightarrow{0}, +^0_R, \overleftarrow{0} \rightarrow +^0_R, \overrightarrow{0} & \overrightarrow{0}, +^1_R, \overleftarrow{0} \rightarrow +^0_R, \overrightarrow{1} \\
\overrightarrow{1}, +^0_R, \overleftarrow{0} \rightarrow +^0_R, \overrightarrow{1} & \overrightarrow{1}, +^1_R, \overleftarrow{0} \rightarrow +^1_R, \overrightarrow{0} \\
\overrightarrow{0}, +^0_R, \overleftarrow{1} \rightarrow +^0_R, \overrightarrow{1} & \overrightarrow{0}, +^1_R, \overleftarrow{1} \rightarrow +^1_R, \overrightarrow{0} \\
\overrightarrow{1}, +^0_R, \overleftarrow{1} \rightarrow +^1_R, \overrightarrow{0} & \overrightarrow{1}, +^1_R, \overleftarrow{1} \rightarrow +^1_R, \overrightarrow{1} \\
\\
\overrightarrow{0}, +^0_R, \overleftarrow{\delta} \rightarrow +^0_R, \overrightarrow{0} & \overrightarrow{0}, +^1_R, \overleftarrow{\delta} \rightarrow +^0_R, \overrightarrow{1} \\
\overrightarrow{1}, +^0_R, \overleftarrow{\delta} \rightarrow +^0_R, \overrightarrow{1} & \overrightarrow{1}, +^1_R, \overleftarrow{\delta} \rightarrow +^1_R, \overrightarrow{0} \\
\\
\overrightarrow{0}, +^0_R \rightarrow +^0_R, \overrightarrow{0} & \overrightarrow{0}, +^1_R \rightarrow +^0_R, \overrightarrow{1} \\
\overrightarrow{1}, +^0_R \rightarrow +^0_R, \overrightarrow{1} & \overrightarrow{1}, +^1_R \rightarrow +^1_R, \overrightarrow{0} \\
\\
\overrightarrow{\delta}, +^0_R, \overleftarrow{\delta} \rightarrow +^0_R, \overrightarrow{\delta} & \overrightarrow{\delta}, +^1_R, \overleftarrow{\delta} \rightarrow +^0_R, \overrightarrow{1}, \overrightarrow{\delta} \\
\overrightarrow{\delta}, +^0_R \rightarrow +^0_R, \overrightarrow{\delta} & \overrightarrow{\delta}, +^1_R \rightarrow +^0_R, \overrightarrow{1}, \overrightarrow{\delta} \\
\\
\overrightarrow{\delta}, \overleftarrow{\epsilon} \rightarrow \overleftarrow{\epsilon}, \overrightarrow{\delta} & \overrightarrow{\delta}, +^0_R \rightarrow \overrightarrow{\delta}
\end{array}$$




Fig. 13. Computing $3 + 1$

B.3 ENUM-SAT

ENUM-SAT is the problem of enumerating all the solutions for an instance of SAT: we want to know *all* the truth assignments of variables for which the formula is satisfiable. We can also consider a particular case of ENUM-SAT: the problem ONESOL-SAT, which consists in returning *only one* valuation satisfying the formula, when the formula is satisfiable.

Reduce stage: $[\text{reduce:enumsat}(x_1 \dots x_n)] = \vec{v}[\text{var}(x_1)] \dots [\text{var}(x_n)] \vec{v}$

$$\begin{array}{cccc} \vec{v}, t \rightarrow \overleftarrow{v}, v & \vec{v}, \overleftarrow{t} \rightarrow t, \vec{v} & \vec{v}, f \rightarrow \overleftarrow{v}_o & \vec{v}_o, \overleftarrow{t} \rightarrow \vec{v}_o \\ \vec{v}, \overleftarrow{v} \rightarrow v & \vec{v}, \overleftarrow{f} \rightarrow f, \vec{v} & \vec{v}_o, \overleftarrow{v} \rightarrow & \vec{v}_o, \overleftarrow{f} \rightarrow \vec{v}_o \end{array}$$

B.4 MAX-SAT

The problem MAX-SAT consists in, given k SAT-formulae, finding the maximum number of formulae satisfiable by the same valuation. This problem is NP-hard and is complete for the class APX — the class of problems approximable in polynomial time with a constant factor of approximation. The problem MAX-SAT can be extended by returning the valuation of variables that satisfies the greater number of formulae among the k ones.

We do not give the corresponding rules for solving MAX-SAT, we just describe the concerned modules. Each formula among the k formulae given in the input is compiled by the same method used previously, and the resulting arrangement of signals for each formula are placed end-to-end. This results in a beam of formulae composed by k sub-beam, one for each formula. The evaluation process is then the same as seen previously.

To compare the number of satisfiable formulae for each valuation, we used the binary adder introduced for #SAT, that we combine with a module comparing two binary number. The reduce phase follows the same idea that for the other variants, except that after comparing two-by-two the number of formulae satisfiable, the greater number is transmitted to the next level of aggregation for the next comparison. Then, at the top of the construction, we can read in the binary representation of the maximal number of satisfiable formulae.

If we also want the truth assignment that satisfies the greater number of formulae, we can easily devise a new module on the basis of the one used for ENUM-SAT.